Data derived weak universal compression

Narayana Santhanam Venkat Anantharam Aleksander Kavcic Wojciech Szpankowski

arXiv:1411.4407v1 [cs.IT] 17 Nov 2014

Abstract

We are motivated by applications that need rich model classes to represent the application, such as the set of all discrete distributions over large, countably infinite supports. But such rich classes may be too complex to admit estimators that converge to the truth with convergence rates that can be uniformly bounded over the entire model class as the sample size increases (uniform consistency). However, these rich classes may still allow for estimators with pointwise guarantees whose performance can be bounded in a model-dependent way. But the pointwise angle has a drawback as well—estimator performance is a function of the very unknown model that is being estimated, and is therefore unknown. Therefore, even if an estimator is consistent, how well it is doing may not be clear no matter what the sample size.

Departing from the uniform/pointwise dichotomy, a new analysis framework is explored by characterizing rich model classes that may only admit pointwise guarantees, yet all information about the unknown model needed to gauge estimator accuracy can be inferred from the sample at hand. To bring focus, we analyze the universal compression problem in this data derived, pointwise consistency framework.

Today, data accumulated in many biological, financial, and other statistical problems stands out not just because of its nature or size, but also because the questions we ask of it are unlike anything we asked before. There is often a tension in these *big data* problems between the need for rich model classes to better represent the application and our ability to handle these classes at all from a mathematical point of view. As an example of why we may want rich model collections, consider the following two examples—the first in prediction and the second in compression.

a) Prediction: Consider insuring the risk of exposure to the Internet as opposed to the simple credit monitoring tools available today. Given the significant number of identity thefts, security breaches, and privacy concerns, insurance of this nature may be highly desirable. How would one model loss here? After all, losses suffered can range from direct loss of property to more intangible, yet very significant damage resulting from lowered credit scores. Designing insurance policies with ceilings on claim payments keeps us in familiar territory mathematically, but also misses the point of why one may want this sort of insurance. We therefore want a *richer* set of candidate loss models that do not impose artificial ceilings on loss.

b) Compression: In compression as well, there have been a lot of attempts to incorporate rich model classes. The approach philosophically closest to the ideas in this paper is hierarchical universal compression proposed in [1]. In hierarchical universal compression one envisions a collection of sub-classes Λ_i, $i \geq 1$. Each subclass Λ_i perhaps represents one aspect of what we know about the source, and is simple enough to allow universal schemes with good guarantees holding uniformly over Λ_i. However, to model the entire problem, these classes are merged to form a richer class $\Lambda = \cup_{i \geq 1} \Lambda_i$. While each Λ_i can be handled with the compression machinery available, the full class Λ may not be.

We have thus run into a fundamental roadblock here. Richness of model classes is often quantified by metrics such as the VC-dimension [2], the Rademacher complexity [3], [4], [5], or the strong compression redundancy [6], [7], [8], [9], [10]. Typically, one looks for estimation algorithms with model-agnostic guarantees based on the sample size—indeed this is the *uniform* consistency dogma that underlies most formulations of engineering applications today. But any such guarantee on estimators on a model class depends on the complexity metrics above—the more complex a class, the worse the guarantees.

The insurance problem or hierarchical universal compression attempt to harness model classes that are too complex to admit estimators with reasonable model-agnostic guarantees (or *uniformly consistent estimators*). Instead the best we can often do is to have guarantees dependent on not just the sample size but on the underlying model

N. Santhanam is with the University of Hawaii, Manoa; nsanthan@hawaii.edu
V. Anantharam is with the University of California, Berkeley; ananth@eecs.berkeley.edu
A. Kavcic is with the University of Hawaii, Manoa; kavcic@hawaii.edu
W. Szpankowski is with Purdue University; spa@cs.purdue.edu

in addition (*pointwise consistent*). This may not be very helpful either—our gauge of how well the estimator is doing is dependent on the very quantity being estimated!

For example in hierarchical universal compression, the full collection Λ may not be strongly compressible as mentioned before. We instead ask for universal schemes with guarantees depending on the source identity—we just require that the excess codelength used by the scheme against any $p \in \Lambda$ matches the redundancy of the specific subclass Λ_j that p belongs to (instead of the redundancy of Λ—which may not be even finite). Such a guarantee then is a model dependent *hierarchical compression* guarantee. This is useful if can we estimate from the data which subclass Λ_i the underlying source belongs to.

Note that even if hierarchical guarantees are possible, it may be impossible using a finite sample to obtain the identity of the subclass the underlying source belongs to. It is possible then that the hierarchical compression guarantee on the redundancy of the underlying source may be just facetious.

We therefore challenge the dichotomy of *uniform* and *pointwise* consistency in the analysis of statistical estimators. Both uniform and pointwise guarantees have their own drawbacks. The former precludes the desired richness of model classes. While the latter allows for rich model classes, it does not provide practical guarantees that can be used in applications.

Instead, we consider a new paradigm positioned in between these two extremes. This framework modifies the world of pointwise consistent estimators—keeping as far as possible the richness of model classes possible but ensuring that all information needed about the unknown model to evaluate estimator accuracy can be gleaned from the data. We call this *data-driven pointwise consistency*.

To bring focus into the theoretical framework, we will formulate and characterize this approach for weak compression over countably infinite alphabets. This approach generalizes our prior work on a related prediction problem, and we compare prediction and compression in the data-driven consistency framework. Going back to hierarchical compression, suppose we turn the question around. How do we characterize, given Λ, *if* it can be decomposed into distinguishable subclasses Λ_i? As it turns out, the answer to the this question follows from our characterization of data-driven weakly compressible classes of distribution.

I. Formulation of problem

Let $\mathcal{P}$ be a collection of distributions over the naturals $\mathbb{N} = \{1, 2, \ldots\}$. Let $\mathcal{P}^\infty$ be the measures induced over infinite sequences of numbers from $\mathbb{N}$ by *i.i.d.* sampling from distributions in $\mathcal{P}$. $\mathcal{P}^\infty$ is called strongly compressible [7], [11] if there is a measure q satisfying

$$\limsup_{n\to\infty} \sup_{p\in\mathcal{P}^\infty} \frac{1}{n} E_p \log \frac{p(X^n)}{q(X^n)} = 0. \tag{1}$$

We can allow for much richer classes if we work with a weaker requirement for the universal measure q. $\mathcal{P}^\infty$ is called weakly compressible [7], [11] if there exists a measure q over infinite sequences of natural numbers such that $\forall p \in \mathcal{P}^\infty$

$$\limsup_{n\to\infty} \frac{1}{n} E_p \log \frac{p(X^n)}{q(X^n)} = 0. \tag{2}$$

Remark Note that both (1) and (2) are usually phrased with encoders or distributions for length n sequences. However, since we will be concerned mainly with the limits, we can use the simpler formulations above. See [12] for a formal explanation of why these formulations are completely equivalent. □

Throughout this paper, the class $\mathcal{P}^\infty$ will be weakly compressible but not necesssarily strongly so. One positive about the weakly universal measure q is that it is essentially as good as the underlying p (in the number of bits used per symbol encoded) for long enough sequences. But the problem is that "long enough" depends on the unknown p since the convergence to limit may not be uniform in (2).

To clarify this point, given any accuracy $\delta > 0$ we ask for an indicator function $\Phi : \mathbb{N}^* \to \{0, 1\}$ that will clarify this point. The function above observes a sequence in $\mathbb{N}^*$, and decides what sequence length is long enough that the normalized KL divergence in (2) above is below δ, and in addition will remain below δ for longer sequences.

From a notational point of view, we require $\Phi(x^i x_{i+1}) \geq \Phi(x^i)$—namely, once Φ indicates that the length is "long enough" that the normalized KL will remain below δ from that point on, it cannot renege later. When Φ turns 1, we say the scheme *enters* the compression game. Furthermore, we require that for all $p \in \mathcal{P}^\infty$,

$$p(\Phi \text{ enters}) = p(X^n : \lim_{n\to\infty} \Phi(X^n) = 1) = 1.$$

Fix a universal measure q. Given $\delta > 0$, Φ is *δ-premature* for a source $p \in \mathcal{P}^\infty$ and string x_1^i if for some $j \leq i$,

$$\Phi(x_1^j) = 1 \text{ and } \frac{1}{j} E_p \log \frac{p(X^j)}{q(X^j)} > \delta.$$

Note that given a measure p, the set of all strings on which Φ is $\delta-$premature can be identified with a prefix free set corresponding to the first times the accuracy condition was violated for the strings. The probability under p of Φ being $\delta-$premature is the probability of this prefix free set.

Definition 1. Given a weakly compressible class $\mathcal{P}^\infty$ we would like to find a universal measure q such that for any accuracy $\delta > 0$ and confidence $\eta > 0$, there is an indicator Φ such that no matter what $p \in \mathcal{P}^\infty$ is in force,

$$p(\Phi \text{ is } \delta-\text{premature}) < \eta.$$

If possible, such a class is weakly compressible in the *data-driven* sense (*d.w.c.*). □

Example 1. Suppose $\mathcal{P}^\infty$ is strongly compressible in addition, namely there exists a measure q satisfying (1). For all $\delta > 0$, the sets

$$N_\delta = \{n : \sup_{p \in \mathcal{P}^\infty} \frac{1}{n} E_p \log \frac{p(X^n)}{q(X^n)} > \delta\}$$

are finite. Suppose we set for any δ, $\Phi(x^i) = 1$ if $i > \max N_\delta$ and 0 else, for all $p \in \mathcal{P}^\infty$ that $p(\Phi \text{ is } \delta-\text{premature}) = 0$. Therefore, the more interesting case is when $\mathcal{P}^\infty$ is weakly compressible, but not strongly compressible. Instead of restricting model classes severely as strong compressibility does, data derived weak compressibility does not depend on the entire class being "simple", but only requires that local neighborhoods be simple. □

In this paper, we obtain a condition that is both necessary and sufficient for an *i.i.d.* class $\mathcal{P}^\infty$ to be data-driven weakly compressible.

Operational justification: The operational justification for our formulation of *d.w.c.* classes of *i.i.d.* sources can be articulated as follows. Given such a class, let q be any measure over infinite length sequences that verifies the definition, i.e. such that for every $\delta > 0$ and $\eta > 0$ there is some $\Phi_{\delta,\eta} : \mathbb{N}^* \mapsto \{0, 1\}$ for which the probability under every p in the model class that $\Phi_{\delta,\eta}$ is δ-premature is less than η.

As we observe the realization of the *i.i.d.* data samples from the (unknown) source p in the model class, we will eventually see a string of some (random) length $n = n(\delta, \eta, p)$ (say x_1^n) such that $\Phi_{\delta,\eta}(x_1^n) = 1$. Now, even though we do not know p, we get the guarantee (with confidence $\geq 1 - \eta$) that using q to compress any subsequent length-n or longer sequence of symbols in the usual way (*i.e.*, $-\log q(x^k)$ bits for a sequence x^k) incurs an expected per-symbol redundancy $\leq \delta$.

I-A. Related formulation for prediction

This framework generalizes a prediction problem a subset of the authors considered in [13]. Suppose we have a collection $\mathcal{P}^\infty$ of *i.i.d.* measures, and samples $X_1, X_2, \ldots$ from an unknown $p \in \mathcal{P}^\infty$. Given a confidence $\eta > 0$, can we come up with a mapping $\Phi : \mathbb{N}^* \to \mathbb{R} \cup \infty$ such that for all p,

$$p(\Phi(X^i) < X_{i+1}) < \eta$$

and Φ is finite eventually almost surely? If so, we say $\mathcal{P}^\infty$ is *insurable*, see [13] for details.

It may not be immediately apparent why insurability and the data derived convergence framework are related. To see the connection, we will frame the insurance problem above a little differently. Let us say we want to consistently bound all percentiles of the underlying distribution using samples from it, in a data-driven way. Denoting the $(1-\delta)$-percentile of a distribution p as $F_p^{-1}(1-\delta)$[1], we find a function $f : \mathbb{N}^* \times [0, 1] \to \mathbb{R}$ and an indicator $I : \mathbb{N}^* \to \{0, 1\}$ as follows. Given a confidence $\eta > 0$, we say we can bound all percentiles consistently in a data-driven fashion, if no matter what $p \in \mathcal{P}$ is in force,

$$p\big(\exists \delta \in (0, 1) : f(X^n, \delta) < F_p^{-1}(1 - \delta) \text{ and } I(X^n) = 1\big) < \eta.$$

[1](will add in appendix for completeness instead of this footnote) see [13] for clarifications regarding what is meant by the inverse cummulative distributive function since the support is discrete here

To clarify the comparisons, note that we call f to be consistent if $\forall p \in \mathcal{P}$,

$$\lim_{n\to\infty} \mathbf{1}\big(\forall \delta \in (0,1), f(X^n,\delta) > F_p^{-1}(1-\delta)\big) = 1 \qquad a.s..$$

Proposition 1. A class $\mathcal{P}^\infty$ is insurable iff we can bound all percentiles consistently in a data derived fashion.
Proof See Appendix I. □

As with compression, we are interested in rich model classes where the rate of convergence to the limit above is not necessarily uniform over the entire model class. In such pointwise convergence cases, the indicator $I(X^n)$ gauges if our estimate f is premature or not using the data—in much the same way as the *d.w.c.* formulation above.

II. Background

II-A. Strong compression

In (1) we defined classes of measures that were strongly compressible. Strong compression redundancy of a class $\mathcal{P}$ of distributions over $\mathbb{N}$ can be seen as the capacity of a channel from $\mathcal{P}$ to $\mathbb{N}$, where the conditional probability distribution over $\mathbb{N}$ given $p \in \mathcal{P}$ is simply the distribution p. The following Lemma simplifies the above viewpoint and yields a simple result we will make use of in the examples to follow.

Lemma 1. Let $\mathcal{P}$ be a class of distributions over a countable support $\mathcal{X}$. For some $m \geq 1$, consider m pairwise disjoint subsets $S_i \subset \mathcal{X}$ ($1 \leq i \leq m$) and let $\delta > 1/2$. If there exist $p_1, \ldots, p_m \in \mathcal{P}$ such that

$$p_i(S_i) \geq \delta,$$

then for all distributions q over $\mathcal{X}$,

$$\sup_{p\in\mathcal{P}} D(p||q) \geq \delta \log m.$$

In particular if there are an infinite number of sets S_i, $i \geq 1$ and distributions $p_i \in \mathcal{P}$ such that $p_i(S_i) \geq \delta$, then the redundancy is infinite.
Proof This is a simplified formulation of the *distinguishability* concept in [14]. For a proof, see *e.g.* [15]. □

We will often be concerned with sequences of symbols drawn *i.i.d.* from $\mathcal{P}$, and let $\mathcal{P}^\infty$ be the measures induced on infinite sequences of naturals obtained by *i.i.d.* sampling from distributions in $\mathcal{P}$. Keeping with convention, we call the redundancy of the class $\mathcal{P}$ to be the single letter redundancy, while the redundancy of $\mathcal{P}^n$ is the length-n redundancy. The question we are concerned about is whether the length-n redundancy, normalized by n, diminishes to 0 as n increases. If so we say that the asymptotic per-symbol redundancy diminishes to 0.

In the *worst case regret* setting, finite single letter regret is necessary and sufficient for the asymptotic per-symbol worst case redundancy to diminish to 0. But in [15], we show that it is not necessarily the case for strong redundancy we are concerned in this paper. It is quite possible that classes with finite single letter strong redundancy have asymptotic per-symbol strong redundancy bounded away from 0.

We therefore need a better handle to analyze the behaviour of length-n strong redundancy in general. To this end, and to avoid future digressions that take away from the thrust of this paper, we have reproduced some relevant results in Appendix III that relate single letter properties to redundancies of length-n sequences.

II-B. Weak compression

Recall that a class $\mathcal{P}^\infty$ of stationary ergodic measures on $\mathbb{N}^\infty$ is defined to be weakly compressible if there is a measure q on $\mathbb{N}^\infty$ that satisfies for all $p \in \mathcal{P}^\infty$ (with finite entropy rate)

$$\lim_{n\to\infty} \frac{1}{n} \mathbb{E}_p \log \frac{p(X^n)}{q(X^n)} = 0,$$

where X^n are sequences of natural numbers from p. Again, it can be shown that the above definition is equivalent to the more commonly used definition from [16], which uses a sequence $q_i : i \geq 1$ of distributions (q_i over length-i sequences) in the left limit. See, *e.g.* [12], for the connection.

In other words, the expected codelength of length-n sequences using the distribution induced by q converges pointwise to the entropy rate over the class $\mathcal{P}^\infty$. Kieffer proved [16] that $\mathcal{P}^\infty$ is weakly compressible iff there

exists a countable set $\mathcal{Q}$ of (single letter) distributions over $\mathbb{N}$ such that for all $p \in \mathcal{P}^\infty$ with finite entropy rate, there exists some distribution $q_p \in \mathcal{Q}$ such that

$$\mathbb{E}_p \log \frac{1}{q_p(X_1)} < \infty,$$

where as before, X_1 is a number chosen from the distribution p. The following corollary of Kieffer's condition will be useful for our proofs.

Corollary 2. If class $\mathcal{P}^\infty$ of measures over $\mathbb{N}^\infty$ is weakly compressible, then there exists a distribution q over $\mathbb{N}$ such that for all $p \in \mathcal{P}^\infty$ with finite entropy rate,

$$\mathbb{E}_p \log \frac{1}{q(X_1)} < \infty.$$

Proof From Kieffer's theorem, we can find a set $\mathcal{Q}$ of single letter distributions over $\mathbb{N}$ such that for all $p \in \mathcal{P}^\infty$ with finite entropy rate and some $q_p \in \mathcal{Q}$

$$\mathbb{E}_p \log \frac{1}{q_p(X_1)} < \infty.$$

Consider the following distribution q over $\mathbb{N}$, that assigns probability

$$q(n) = \sum_{i=1}^{|\mathcal{Q}|} \frac{q_i(n)}{i(i+1)}$$

where the upper limit of summation is understood to be ∞ if $\mathcal{Q}$ is countably infinite. The corollary follows by noting that for all i and for all n,

$$q(n) \geq \frac{q_i(n)}{i(i+1)}.$$

□

III. Summary of main results

Our primary results are necessary and sufficient conditions for *d.w.c.*. We follow it up by examining connections between various compression formulations (strong, *d.w.c.* and weak) with insurability as defined in Section I.

Very complex local neighborhoods kill *d.w.c.*. An indicator scheme could be "deceived" by some process $p \in \mathcal{P}^\infty$ into certifying accuracy, while a close enough distribution lurks with very bad performance. Note that since $\mathcal{P}^\infty$ are *i.i.d.*, the sources therein can be identified without confusion using their single letter marginals as well.

For any two measures p and q, we let

$$D_n(p||q) \stackrel{\text{def}}{=} E_{p(X^n)} \log \frac{p(X^n)}{q(X^n)},$$

the KL divergence between the distributions induced over length n sequences by p and q respectively. Furthermore, for measures p and q,

$$\mathcal{J}(p,q) = D_1\left(p||\frac{p+q}{2}\right) + D_1\left(q||\frac{p+q}{2}\right),$$

where in the above, the KL divergences are taken between the single letter distributions corresponding to p and q.

An $\epsilon-$*neighborhood* of $p \in \mathcal{P}^\infty$ is the set $B(p,\epsilon)$ of all sources $p' \in \mathcal{P}^\infty$ such that $\mathcal{J}(p,p') < \epsilon$. From Lemma 11 in the Appendix, note that $B(p,\epsilon)$ is completely contained in an ℓ_1-neighborhood of radius $\epsilon \ln 2$.

III-A. Deceptive measures

Roughly speaking, $p \in \mathcal{P}^\infty$ is *deceptive* if the strong redundancy of neighborhoods of p is bounded away from 0 in the limit as the neighborhood shrinks to 0. More precisely, we say p is deceptive if

$$\lim_{\epsilon \to 0} \inf_{q} \limsup_{n \to \infty} \sup_{p' \in B(p,\epsilon)} \frac{1}{n} D_n(p'||q) > 0.$$

In Section IV we show several examples of classes with deceptive sources, as well as classes with no deceptive sources.

Lemma 3. If $p \in \mathcal{P}^\infty$ is not deceptive, then there is a measure q^* such that

$$\lim_{\epsilon \to 0} \limsup_{n \to \infty} \sup_{p' \in B(p,\epsilon)} \frac{1}{n} D_n(p'||q^*) = 0.$$

Proof Consider the sequence of neigborhoods $B(p, 1/m)$. Because p is not deceptive, we can find a sequence of measures q_m such that

$$\limsup_{n \to \infty} \sup_{p' \in B(p,1/m)} \frac{1}{n} D_n(p'||q_m) \leq \delta_m + \frac{1}{m},$$

where $\lim_{m \to \infty} \delta_m \to 0$. The lemma is satisfied by the measure that assigns probability to any sequence $\mathbf{x}$

$$q^*(\mathbf{x}) = \sum_{m \geq 1} \frac{q_m(\mathbf{x})}{m(m+1)}.$$ □

III-B. Main result

Our main result relates deceptive measures to *d.w.c.*.

Theorem 4. $\mathcal{P}^\infty$ is *d.w.c.* iff no $p \in \mathcal{P}^\infty$ is deceptive. □

Remark We show in [13] that $\mathcal{P}^\infty$ is insurable iff some neighborhood (as defined here) of every $p \in \mathcal{P}^\infty$ is tight[2]. In both cases, note that the condition relies on some neighborhood of every model being simple. This appears to be a feature of all such data-driven formulations, see the Conclusion as well. From Lemma 8, we therefore obtain

Corollary 5. If $\mathcal{P}^\infty$ is *d.w.c.*, then $\mathcal{P}^\infty$ is insurable. □

III-C. Connections between compression and prediction

As one would expect, we show that strong compression implies insurability. But the relation between weak compression and insurability is more subtle—we show that neither implies the other. We construct classes that are weakly compressible but not insurable, as well as those that are insurable but not weakly compressible. Finally, as shown in the Corollary above, *d.w.c.* implies insurability.

IV. Examples

We construct a series of examples that highlight various aspects of our formulation as well as that of deceptive distributions.

The first collection we consider is $\mathcal{U}$, the collection of all uniform distributions over finite supports of form $\{m, m+1, \ldots, M\}$ for all positive integers m and M with $m \leq M$. Let the sequence of losses be *i.i.d.* samples from distributions in $\mathcal{U}$—call the resulting model class over infinite loss sequences $\mathcal{U}^\infty$.

The second collection is the set $\mathcal{N}_1^\infty$ of all *i.i.d.* processes such that the one dimensional marginals have finite first moment. Namely, $\forall p \in \mathcal{N}_1^\infty$, $E_p X < \infty$ where $X \in \mathbb{N}$ is distributed according to the single letter marginal of p. Let $\mathcal{N}_1$ be the collection of single letter marginals from $\mathcal{N}_1^\infty$. It is easy to that $\mathcal{U} \subseteq \mathcal{N}$. It is easy to verify as below that every distribution in $\mathcal{N}_1$ is deceptive.

[2]To be defined

A monotone probability distribution p on $\mathbb{N}$ is one that satisfies $p(y+1) \leq p(y)$ for all $y \in \mathbb{N}$. We will also consider $\mathcal{M}$, the collection of monotone distributions on $\mathbb{N}$ with finite entropy. Let $\mathcal{M}^\infty$ be the set of all *i.i.d.* processes, with one dimensional marginals from $\mathcal{M}$.

Example 2. $\mathcal{U}^\infty$ is not strongly compressible but is *d.w.c.*.

Proof We first show that $\mathcal{U}$ has infinite redundancy. To see this, we partition $\mathbb{N}$ into disjoint T_i, $i \geq 0$, where $T_i = \{2^i, \ldots, 2^{i+1}-1\}$ be a set of 2^i elements. We will use this partitioning of $\mathbb{N}$ multiple times in the paper.

Furthermore for each T_i there is an associated distribution $p_i \in \mathcal{U}$ such that $p_i(T_i) = 1$. From Lemma 1, we therefore obtain the redundancy of $\mathcal{U}$ to be ∞. Therefore the redundancy of $\mathcal{U}^n$ is infinite for all n.

To see the class is *d.w.c.*, observe that around each distribution in $\mathcal{U}$ is an ℓ_1-neighborhood that contains no other distribution of $\mathcal{U}$. □

Example 3. $\mathcal{N}_1^\infty$ is weakly compressible but not *d.w.c.*.

Proof By definition, all $p \in \mathcal{N}_1$ satisfy $\sum_{i\geq 1} ip_i < \infty$. Therefore Kieffer's condition in Corollary 2 is satisfied by the distribution $q(i) = 1/2^i$ ($i \geq 1$). Thus we observe that $\mathcal{N}_1$ is weakly compressible.

We show every distribution in $\mathcal{N}_1$ is deceptive by showing that no neighborhood around any distribution in $\mathcal{N}_1$ is tight. From Lemma 8, we then conclude that no neighborhood around any distribution in $\mathcal{N}_1$ has finite strong redundancy. Therefore, the asymptotic per-symbol redundancy of *i.i.d.* sequences is not finite either, and $\mathcal{N}_1$ is not *d.w.c.*.

To see that no neighborhood of any $p \in \mathcal{N}_1$ is tight, consider distributions of form $p' = (1-\epsilon)p + \epsilon q$, where $q \in \mathcal{U}$ is a monotone uniform distribution and $\epsilon > 0$. Since q has finite support, $p' \in \mathcal{N}_1$ as well. Clearly, the ℓ_1 distance between p' and q is $\leq 2\epsilon$. But for all $M > 0$ and $\delta < \epsilon$, we can pick $q \in \mathcal{U}$ over a sufficiently large support that the $1-\delta-$percentile of p' can be made $\geq M$. The observation follows. □

Example 4. $\mathcal{M}^\infty$ is weakly compressible but not *d.w.c.*.

Proof We first observe using Kieffer's condition that $\mathcal{M}^\infty$ is weakly compressible. For all $p \in \mathcal{M}$ and all numbers n, we have

$$p(n) \leq \frac{1}{n}.$$

It follows that every $p \in \mathcal{M}$ with finite entropy must satisfy

$$\sum_{n\geq 1} p(n) \log n \leq \sum_{n\geq 1} p(n) \log \frac{1}{p(n)} < \infty.$$

Now consider the distribution q over $\mathbb{N}$ assigning probabilities $q(n) = \frac{6}{\pi^2 n^2}$. The equation above now implies that for all $p \in \mathcal{M}$ with finite entropy,

$$\sum_{n\geq 1} p(n) \log \frac{1}{q(n)} < \infty.$$

To see that every distribution $p \in \mathcal{M}$ is deceptive, we proceed like in the previous example and show that no neighborhood around p is tight. Therefore, no neighborhood around p can have finite single letter redundancy (and hence the asymptotic per-symbol redundancy cannot be finite either). Consider distributions of form $p' = (1-\epsilon)p + \epsilon q$, where $q \in \mathcal{U} \cap \mathcal{M}$ is a monotone uniform distribution and $\epsilon > 0$.

The ℓ_1 distance between p' and q is $\leq 2\epsilon$. For all $M > 0$ and $\delta < \epsilon$, we can pick $q \in \mathcal{U}$ over a sufficiently large support that the $1-\delta-$percentile of p' can be made $\geq M$. Thus, no neighborhood around p' is tight, and consequently no neighborhood is compressible either from Lemma 8. □

Now for $h > 0$, we consider the set $\mathcal{M}_h \subset \mathcal{M}$ of all monotone distributions over $\mathbb{N}$ such that the second moment of the self information,

$$E_p\left(\log \frac{1}{p(X)}\right)^2$$

is bounded above by h. Let $\mathcal{M}_h^\infty$ be the set of all *i.i.d.* loss processes with one dimensional marginals from $\mathcal{M}_h$. Then

Example 5. $\mathcal{M}_h^\infty$ is strongly compressible, hence *d.w.c.*.

Proof Note that for any monotone p and all naturals i, $p(i) \leq 1/i$. Therefore for any $p \in \mathcal{M}_h$, if $X \sim p$ we have

$$E_p \log^2(X+1) \leq E_p \log^2 \frac{1}{p(X)} \leq h.$$

Hence, we have for the distribution q over $\mathbb{N}$ assigning $q(i) = \frac{1}{i(i+1)}$,

$$E_p\left(\log \frac{p(X)}{q(X)}\right)^2 \leq E_p(\log(X+1))^2 \leq h.$$

From [15] (reproduced for this special case in Appendix III), we can therefore construct a measure q^* such that

$$\sup_{p \in \mathcal{M}_h^\infty} \frac{1}{n} D_n(p||q) \leq \frac{2h^{3/2}}{\sqrt{\log n}} + \pi\sqrt{\frac{2}{3n}}$$

Namely the class $\mathcal{M}_h^\infty$ is strongly compressible, and therefore *d.w.c.* trivially from Example 1. □

In the class $\mathcal{U}$ above, there was a neighborhood around each distribution $p \in \mathcal{U}$ with no other model from $\mathcal{U}$. Hence $\mathcal{U}$ trivially satisfied the local redundancy condition of Theorem 4. The $\mathcal{M}_h$ case falls into another extreme—the entire model class $\mathcal{M}_h$ is strongly compressible, and therefore the conditions of Theorem 4 was satisfied in a trivial way again. The following example illustrates a *d.w.c.* class of models where neither extreme holds.

For a distribution q over $\mathbb{N}$, let $q^{(R)}(i+R) = q(i)$ for all $i \in \mathbb{N}$. Furthermore let the *span* of any finite support probability distribution over naturals be the largest natural number which has non-zero probability. Then, let

$$\mathcal{F}_h = \left\{(1-\epsilon)p_1 + \epsilon p_2^{(\text{span}(p_1)+1)} : \forall p_1 \in \mathcal{U}, p_2 \in \mathcal{M}_h \text{ and } 1 > \epsilon > 0\right\}$$

As always $\mathcal{F}_h^\infty$ is the set of measures on infinite sequences formed by *i.i.d.* sampling from distributions in $\mathcal{F}_h$.

Example 6. $\mathcal{F}_h^\infty$ is *d.w.c.*.

Proof Let the *base* of any probability distribution over the naturals be the smallest natural number which has non-zero probability. Consider any distribution $p = (1-\epsilon)p_1 + \epsilon p_2^{(\text{span}(p_1)+1)} \in \mathcal{F}_h$ with $p_1 \in \mathcal{U}$, $p_2 \in \mathcal{M}_h$, and $1 > \epsilon > 0$. Let m denote base(p), and $m + M - 1$ denote span(p_1). Thus $|\text{support}(p_1)| = M$, and we have $M \geq 1$. Of course, we also have base(p_1) = base(p).

Consider any other distribution $q = (1-\epsilon')q_1 + \epsilon' q_2^{(\text{span}(q_1)+1)} \in \mathcal{F}_h$, where $q_1 \in \mathcal{U}$, $q_2 \in \mathcal{M}_h$, and $1 > \epsilon' > 0$. If the

$$|\text{support}(q_1)| + \text{base}(q_1) > m + \left\lceil \frac{M}{1-\epsilon} \right\rceil,$$

we have that there is at least one element in the support of p_1 that contributes

$$\frac{(1-\epsilon)}{M} - \frac{1}{\frac{M}{1-\epsilon}+1} = \frac{(1-\epsilon)^2}{M(M+1)}$$

to the ℓ_1 distance between p and q.

Consider the set of all distributions with base m' and support of the first component in $\mathcal{U}$ equal to M'. This set of distributions is clearly strongly compressible using the argument in Example 5. Furthermore note that a finite union of strongly compressible sets of distributions is also strongly compressible.

Note that the set of all distributions q within a ℓ_1 ball of radius $\frac{(1-\epsilon)^2}{M(M+1)}$ satisfy $|\text{support}(q_1)| + \text{base}(q_1) \leq \left\lceil \frac{M}{1-\epsilon} \right\rceil$, and can hence be decomposed into a finite union of strongly compressible collections of distributions. Therefore the set of distributions in $\mathcal{F}_h$ within ℓ_1 distance from p is strongly compressible. Thus, no distribution is deceptive and the class is *d.w.c.*. □

V. Necessary part

This side of the characterization follows very naturally from the definition of deceptive measures. We follow the characterization with qualifications on how to interpret the result.

Theorem 6. $\mathcal{P}^\infty$ is *d.w.c.* only if no $p \in \mathcal{P}^\infty$ is deceptive.

Proof We prove the contrapositive of the statement above. Namely we show that if some $p \in \mathcal{P}^\infty$ is deceptive, then $\exists \eta > 0$ and $\delta > 0$ such that $\forall$ indicator schemes Φ, there is some $p' \in \mathcal{P}^\infty$ such that $p'(\Phi \text{ is } \delta-\text{premature}) > \eta$.

To pick η, choose any $\alpha > 0$, and pick $\eta = 1 - \alpha$. Since p is deceptive, we can pick a δ that is > 0 and

$$< \lim_{\epsilon \to 0} \limsup_{n \to \infty} \sup_{p' \in B(p,\epsilon)} \frac{1}{n} D_n(p'||q).$$

The rest of the proof applies for all measures q and all indicator schemes Φ. For all $n \geq 1$, let

$$R_n \stackrel{\text{def}}{=} \{x^n : \Phi(x^n) = 1\}$$

be the set of sequences of length n on which Φ has entered and let $N \geq 4/\alpha$ be a number such that $p(R_N) > 1 - \alpha/2$.

Set[3] $\epsilon = \frac{1}{16(\ln 2)N^8}$. Applying Lemma 12 to distributions over length-N sequences induced by p and any $\tilde{p} \in \mathcal{P}^\infty$ such that $\mathcal{J}(p, \tilde{p}) \leq \epsilon$, we have

$$\tilde{p}(R_N) \geq 1 - \alpha/2 - \frac{2}{N} \geq 1 - \alpha. \tag{3}$$

Note that

$$\inf_q \limsup_{n \to \infty} \sup_{p' \in B(p,\epsilon)} \frac{1}{n} D_n(p'||q)$$

is non-increasing with ϵ, and that the limit as $\epsilon \to 0$ is $> \delta$. Therefore, we can choose $n > N$ and $\tilde{p} \in B(p, \epsilon)$ such that

$$\tilde{p}(R_n) \geq 1 - \alpha \text{ and } \frac{1}{n} D_n(\tilde{p}||q) > \delta.$$

This in turn means for the choice of η and δ above, $\tilde{p}(\Phi \text{ is } \delta-\text{premature }) > \eta$. Because Φ and q were arbitrary, the theorem follows. □

When a class $\mathcal{P}^\infty$ is not *d.w.c.*, all we know is that the class has at least one deceptive measure. If a measure is deceptive, it does not automatically imply that any other measure in any neighborhood (no matter how close to the deceptive measure) is deceptive. It is therefore quite possible that a class has only a few deceptive distributions, and removing the deceptive measures would make the class *d.w.c.*. The following example illustrates such a case.

Example 7. As before, we partition the set $\mathbb{N}$ into $T_i = \{2^i, \ldots, 2^{i+1} - 1\}$, $i \geq 0$. Note that T_i has 2^i elements. For $\epsilon > 0$, let $n_\epsilon = \lfloor \frac{1}{\epsilon} \rfloor$. Let $1 \leq j \leq 2^{n_\epsilon}$ and let $p_{\epsilon,j}$ be a distribution on $\mathbb{N}$ that assigns probability $1 - \epsilon$ to the number 1 (or equivalently, to the set T_0), and ϵ to the j'th smallest element of T_{n_ϵ}, namely the number $2^{n_\epsilon} + j - 1$. Furthermore, let p_0 be a singleton distribution assigning probability 1 to the number 1. $\mathcal{B}$ (mnemonic for binary, since every distribution has at most support of size 2) is the collection of distributions $p_{\epsilon,j}$ for all $\epsilon > 0$ and p_0. $\mathcal{B}^\infty$ is the set of measures over infinite sequences of numbers corresponding to *i.i.d.* sampling from $\mathcal{B}$.

From Lemma 8, we obtain that for any neighborhood of p_0, $B(p_0, \delta)$, and any measure q over infinite sequences of numbers (see [15] for details)

$$\limsup_{n \to \infty} \sup_{p \in B_{p_0,\delta}} \frac{1}{n} D_n(p||q) \geq \frac{1}{e}.$$

Therefore, p_0 is deceptive. On the other hand, around every other distribution $p_{\epsilon,j} \in \mathcal{P}$ is a ball of ℓ_1 distance $\frac{2}{n_\epsilon}$ that contains only distributions with the same support namely $\{1, 2^{n_\epsilon} + j - 1\}$. Thus,

$$\limsup_{n \to \infty} \sup_{p : |p - p_{\epsilon,j}| \leq \frac{2}{n_\epsilon}} \frac{1}{n} D_n(p||q) \leq \lim_{n \to \infty} \frac{\frac{1}{2} \log n}{n} = 0.$$

[3]Please note that in the interest of simplicity, we have not attempted to provide the best scaling for ϵ or the tightest possible bounds in arguments below

Essentially, this implies that p_0 is deceptive but no other measure is. Indeed, this is quite clear when we go about compressing sequences from $\mathcal{B}^\infty$. If at any point we see two symbols in the string, from that point on there is no ambiguity in the support and very little ambiguity in probabilities.

But if we see a sequence of all 1s, we can never be sure (no matter what the length of sequence of 1s) if the source is indeed p_0. Given m ones, we do not know if the source may be $p_{\epsilon,j}$ for some $\epsilon \ll \frac{1}{m}$ and $j \in T_{n_\epsilon}$. And if indeed the source is such a $p_{\epsilon,j}$, it is impossible to insure good performance for the yet unseen symbol—if we have seen a sequence of m ones, the codelength loss when we first unseen symbol is revealed is $\gg 2^m$. This holds no matter how large m is.

However, if we remove p_0 from the class, we have no such trouble. We have no obligation to have a starting point on the sequence on infinite 1s since this sequence has probability 0 under every remaining source. □

Countable unions of compressible sources

One should guard against assuming that any countable collection of compressible classes is *d.w.c.*. In fact, the class $\mathcal{B}^\infty$ in Example 7 is again a counterexample to such a notion. It is necessary that every source in the union is non-deceptive. The fact that every source was non-deceptive in the subclass it belonged to is not a guarantee that the source remains non-deceptive in the union, as with $p_0 \in \mathcal{B}$ above.

A similar example can be constructed from $\mathcal{M}_h$ in Example 5. Though $\mathcal{M}_h$ is strongly compressible for every $h > 0$, essentially the same argument as in Example 4 proves that the union $\cup_{h\in\mathbb{N}}\mathcal{M}_h$ is not *d.w.c.*. To see this, observe that all the distributions that are considered in Example 4 also have their second moment of self-information finite.

With reference to hierarchical compression, suppose $\Lambda = \cup_{i\geq 1}\Lambda_i$ is a disjoint countable union of strongly compressible classes Λ_i. While hierarchical compression schemes can ensure that the excess codelength of the scheme against any $p \in \Lambda$ matches the redundancy of the subclass Λ_i that p belongs to, the question we consider is whether we can find out the identity of the class Λ_i from the data.

We will argue that every source in Λ has to be non-deceptive if we are to figure out which Λ_i the underlying source lies almost surely in finite[4] time. To see this, first note that no source is deceptive in any single subclass Λ_i. If a deceptive source does exist in Λ, there must hence be sources from other subclasses in any arbitrarily small neighborhood. Furthermore, no neighborhood of a deceptive source in Λ can have sources from finitely many other neighborhoods since finite collections of strongly compressible sources are automatically strongly compressible. Therefore deceptive sources must have in any neighborhood, however small, sources from an uncountably infinite number of subclasses.

If there is a deceptive source in Λ, we hence cannot pinpoint the membership of the source into a subclass Λ_i in a data-driven sense.

On the contrary, suppose we are given a class Λ where no source is deceptive. Can we obtain a countable union $\Lambda = \cup_{i\geq 1}\Lambda_i$ so that we can establish the membership of the underlying source (almost surely in finite time, no matter what it is) in one of the Λ_i the source lies in? The answer turns out to be positive from the result on sufficiency for *d.w.c.* in the next Section.

VI. Sufficient part

When no $p \in \mathcal{P}^\infty$ is deceptive, we construct a measure q such that given any confidence $\eta > 0$ and accuracy δ, there is a indicator scheme Φ that is $\delta-$premature with probability $\leq \eta$.

From Lemma 3, if no $p \in \mathcal{P}^\infty$ is deceptive, there is for each $p \in \mathcal{P}^\infty$ a neighborhood $B(p, \epsilon_p)$ such that

$$\limsup_{n\to\infty} \sup_{p'\in B(p,\epsilon_p)} \frac{1}{n} D_n(p'||q^*) < \delta.$$

We pick such a neighborhood $B(p, \epsilon_p)$ for each $p \in \mathcal{P}$ and call it the *reach* of p. The reach of p will play the role of the set of measures in $\mathcal{P}^\infty$ for which it will be okay to eventually set indicators assuming p is in force.

[4]note that the time need not be bounded over all sources

VI-A. Topology of $\mathcal{P}$ with the ℓ_1 metric

To prove that $\mathcal{P}^\infty$ is *d.w.c.* if no measure is deceptive, we will need to find a way to cover $\mathcal{P}$ with countably many sets of the form $B(p, \epsilon_p)$ above. Unfortunately, $\mathcal{J}(p, q)$ is not a metric, so it is not immediately clear how to go about doing this. On the other hand note that

$$\mathcal{J}(p', p) \leq |p - p'|_1 / \ln 2,$$

where $|p - p'|_1$ denotes the ℓ_1 distance between the single letter marginals of p and p' (see Lemma 11 in the Appendix). Therefore, we can instead bootstrap off an understanding of the topology induced on $\mathcal{P}$ by the ℓ_1 metric.

The topology induced on $\mathcal{P}$ by the ℓ_1 metric is Lindelöf, i.e. any covering of $\mathcal{P}$ with open sets in the ℓ_1 topology has a countable subcover (see [17, Defn. 6.4] for definitions and properties of Lindelöf topological spaces). See [13] for the proof of why $\mathcal{P}$ is Lindelöf.

VI-B. Sufficient condition

We now have the machinery required to prove that if no $p \in \mathcal{P}^\infty$ is deceptive, then $\mathcal{P}^\infty$ is *d.w.c.*.

Theorem 7. If no $p \in \mathcal{P}$ is deceptive, then $\mathcal{P}^\infty$ is *d.w.c.*.

Proof The proof is constructive. For any confidence $0 < \eta < 1$ and accuracy δ, we obtain an indicator scheme Φ such that for all $p \in \mathcal{P}^\infty$,

$$p(\Phi \text{ is } \delta-\text{premature }) < \eta.$$

Wherever we use ℓ_1 distances $|p - \tau|_1$, it will be understood that we mean the one dimensional marginals of the measures p and τ respectively.

For $p \in \mathcal{P}$, define the following set

$$Q_p = \left\{ \tau : |p - \tau|_1 < \frac{{\epsilon_p}^2 (\ln 2)^2}{16} \right\},$$

where ϵ_p is the reach of p, and τ above is any distribution over $\mathbb{N}$ (not necessarily in $\mathcal{P}$). We will call Q_p as the *zone* of p. The set Q_p is non-empty when $\epsilon_p > 0$.

For large enough n, the set of sequences of length n with empirical distribution in Q_p will ensure that the indicator scheme Φ to be proposed enters with probability 1 when p is in force. Note that if $\epsilon_p > 0$ is small enough then $Q_p \cap \mathcal{P} \subset B(p, \epsilon_p)$—we will assume wolog that $\epsilon_p > 0$ is always taken so that $Q_p \cap \mathcal{P} \subset B(p, \epsilon_p)$.

Since no $p \in \mathcal{P}$ is deceptive, none of the zones Q_p are empty and trivially $p \in Q_p \cap \mathcal{P}$. Thus, the space $\mathcal{P}$ of distributions can be covered by the sets $Q_p \cap \mathcal{P}$, namely

$$\mathcal{P} = \cup_{p \in \mathcal{P}} (Q_p \cap \mathcal{P}).$$

From Section VI-A, we know that $\mathcal{P}$ is Lindelöf under the ℓ_1 topology. Thus, there is a countable set $\tilde{\mathcal{P}} \subseteq \mathcal{P}$, such that $\mathcal{P}$ is covered by the collection of relatively open sets

$$\{ Q_{\tilde{p}} \cap \mathcal{P} : \tilde{p} \in \tilde{\mathcal{P}} \}.$$

We let the above collection be denoted by $\mathcal{Q}_{\tilde{\mathcal{P}}}$. We will refer to $\tilde{\mathcal{P}}$ as the *quantization* of $\mathcal{P}$ and to elements of $\tilde{\mathcal{P}}$ as *centroids* of the quantization, borrowing from commonly used literature in classification.

We index the countable set of centroids, $\tilde{\mathcal{P}}$ (and reuse the index for the corresponding elements of $\mathcal{Q}_{\tilde{\mathcal{P}}}$) by $\iota : \tilde{\mathcal{P}} \to \mathbb{N}$.

c) Description of q^*: For each $\tilde{p} \in \tilde{\mathcal{P}}$, from Lemma 3 we have a locally universal measure $\tilde{q}$ such that

$$\limsup_{n\to\infty} \sup_{p' \in B(\tilde{p}, \epsilon_{\tilde{p}})} \frac{1}{n} D_n(p'||\tilde{q}) < \delta.$$

Let $\tilde{R}$ be the collection of these locally universal measures for all $\tilde{p} \in \tilde{\mathcal{P}}$. Note that $\tilde{R}$ is countable and set the index $\iota(\tilde{q})$ to be equal the index assigned to its the corresponding centroid $\tilde{p}$ in the enumeration of $\tilde{\mathcal{P}}$. Then for all sequences $\mathbf{x}$

$$q^*(\mathbf{x}) = \sum_{\tilde{q} \in \tilde{R}} \frac{\tilde{q}(\mathbf{x})}{\iota(\tilde{q})(\iota(\tilde{q})+1)}$$

Observe again from Lemma 3 and the above quantization that for all $p \in \mathcal{P}^\infty$,

$$\limsup_{n\to\infty} \frac{1}{n} D_n(p||q^*) < \delta.$$

Moreover for all $\tilde{p} \in \tilde{\mathcal{P}}$,

$$\limsup_{n\to\infty} \sup_{p' \in B(\tilde{p}, \epsilon_{\tilde{p}})} \frac{1}{n} D_n(p||q^*) < \delta. \tag{4}$$

We now construct the indicator scheme Φ having the property that for all $p \in \mathcal{P}^\infty$,

$$p(\Phi \text{ is } \delta-\text{premature }) < \eta.$$

d) Preliminaries: Consider a length-n sequence x^n on which Φ has not entered thus far. Let the empirical distribution of the sequence be q, and let

$$\mathcal{P}'_\tau := \{p' \in \tilde{\mathcal{P}} : \tau \in Q_{p'}\}$$

be the set of centroids in the quantization of $\mathcal{P}$ (elements of $\tilde{\mathcal{P}}$) which can potentially *capture* τ. Note that τ in general need not belong to $\tilde{\mathcal{P}}$ or $\mathcal{P}$.

If $\mathcal{P}'_\tau \neq \emptyset$, we will further refine the set of distributions that could capture τ further to $\mathcal{P}_\tau \subset \mathcal{P}'_\tau$ as described below. Refining $\mathcal{P}'_\tau$ to $\mathcal{P}_\tau$ ensures that models in $\mathcal{P}'_\tau$ do not $\delta-$prematurely capture sequences.

Let p be the model in force, which remains unknown. The idea is that we want sequences generated by (unknown) p to be captured by those centroids of the quantization $\tilde{\mathcal{P}}$ that have p in their reach. We will require (5) below to ensure that the probability (under the unknown p) of all sequences that may get captured by centroids $p' \in \mathcal{P}_\tau$ not having p in its reach remains small. In addition, we impose (6) as well to resolve a technical issue since τ need not, in general, belong to $\mathcal{P}$.

For $p' \in \mathcal{P}'_\tau$, let the reach of p' be $\epsilon_{p'}$, and define

$$D_{p'} := \frac{{\epsilon_{p'}}^4 (\ln 2)^4}{256} .$$

In case the underlying distribution p happens to be out of the reach of p' (wrong capture), the quantity $D_{p'}$ will later lower bound the distance of the empirical τ in question from the underlying p.

Specifically, we place p' in $\mathcal{P}_\tau$ if n satisfies

$$\exp\left(-nD_{p'}/18\right) \leq \frac{\eta}{2C(p')\iota(p')^2 n(n+1)}, \tag{5}$$

and

$$2F_\tau^{-1}(1-\sqrt{D_{p'}}/6) \leq \log C(p'), \tag{6}$$

where for any $0 < \gamma < 1$, $F_\tau^{-1}(1-\gamma)$ is the $1-\gamma$ percentile of τ as defined in [13]. where $C(p')$ is

$$C(p') := 2^{2^{\left(\sup_{r \in B(p', \epsilon_{p'})} F_r^{-1}(1-\sqrt{D_{p'}}/6)\right)}}.$$

Note that $C(p')$ is finite from Lemma 8 and because p' is not deceptive. See [13] for why the above equations look this way.

e) Description of Φ*:* For the sequence x^m with type τ, if $\mathcal{P}_\tau = \emptyset$ the scheme does not enter yet. If $\mathcal{P}_\tau \neq \emptyset$, let p_τ denote the distribution in $\mathcal{P}_\tau$ with the smallest index. All sequences with prefix x^m are then said to be *trapped* by p_τ.

From (4),

$$\limsup_{n\to\infty} \sup_{p'\in B(p_q,\epsilon_{p_q})} \frac{1}{n} D_n(p||q^*) < \delta,$$

therefore the set

$$N_{p_\tau} = \{n : \sup_{p'\in B(p_q,\epsilon_{p_q})} \frac{1}{n} D_n(p||q^*) \geq \delta\}$$

is finite. If $m > \max N_{p_\tau}$, we set $\Phi(x^m) = 1$, 0 else.

f) Φ *enters with probability 1:* First, we verify that the scheme is trapped with probability 1, no matter what distribution $p \in \mathcal{P}$ is in force. From the previous paragraph, this also means that Φ enters with probability 1.

To see that the scheme is trapped with probability 1 no matter which $p \in \mathcal{P}$ is in force, please see an identical argument in [13]. For completeness, we reproduce this in Appendix IV.

g) Probability Φ $\delta-$*premature* $\leq \eta$*:* We now analyze the scheme. Consider any $p \in \mathcal{P}$. Among sequences on which Φ has entered, we will distinguish between those that are in *good* traps and those in *bad* traps. If a sequence x^n is trapped by p' such that $p \in B(p', \epsilon_{p'})$, p' is a good trap. Conversely, if $p \notin B(p', \epsilon_{p'})$, p' is a bad trap.

(Good traps) Suppose a length-n sequence x^n is in a good trap, namely, it is trapped by a distribution p' such that $p \in B(p', \epsilon_{p'})$. In this case, we therefore have

$$p(\Phi \text{ is } \delta-\text{premature}) = 0.$$

(Bad traps) We can show that the probability with which sequences generated by p fall into bad traps $\leq \eta$ using an argument identical to [13], reproduced in Appendix V for completeness sake. Pessimistically, we assume that Φ is $\delta-$premature on every sequence that falls into a bad trap.

The theorem follows. □

VII. Connections with prediction

As mentioned before, the prediction problem—or insurability as in [13] appears closely connected to our formulation of the *d.w.c.* here. A class of *i.i.d.* measures $\mathcal{P}^\infty$ is insurable iff some neighborhood of every distribution among the single letter marginals ($\mathcal{P}$) is tight.

We now examine connections between the insurance problem on the one hand, and various formulations of universal compression on the other. We show that if $\mathcal{P}$ can be strongly compressed, it can be insured as well. However, the connection with weak compression is more subtle. We show by constructing appropriate classes of distributions that neither weak compression nor insurability implies the other. Finally, *d.w.c.* implies insurability. But not all insurable classes are even weakly compressible, leave alone *d.w.c.*.

VII-A. Strong compression and insurability

We first show that strong compression implies insurability. On the other hand, there can be insurable classes that are not strongly compressible as seen from Example 2.

Lemma 8. A class $\mathcal{P}$ with bounded strong redundancy is tight. Namely, if the strong redundancy of $\mathcal{P}$ is finite, then for any $\gamma > 0$

$$\sup_{p\in\mathcal{P}} F_p^{-1}(1-\gamma) < \infty.$$

Proof $\mathcal{P}$ has bounded strong redundancy. Let q be a distribution over $\mathbb{N}$ such that

$$\sup_{p\in\mathcal{P}} D(p||q) < \infty,$$

and we define $R = \sup_{p\in\mathcal{P}} D(p||q)$. It follows that for all $p \in \mathcal{P}$ and any m,

$$p(\left|\log \frac{p(X)}{q(X)}\right| > m) \le (R + (2\log e)/e)/m,$$

To see the above, note that if S is the set of all numbers such that $p(x) < q(x)$, a well-known convexity argument shows that

$$\sum_x p(x)\log\frac{p(x)}{q(x)} \ge p(S)\log\frac{p(S)}{q(S)} \ge -\frac{\log e}{e}.$$

We prove the lemma by contradiction. Pick m so large that $(R + (2\log e)/e)/m < \gamma/2$. For all p, we show that

$$p\big(x : x \ge F_q^{-1}(1-\gamma/2^{m+1})\big) \le \gamma.$$

To see the above, observe that we can split the tail $x \ge F_q^{-1}(1-\gamma/2^{m+1})$ into two parts—(i) numbers x such that $\log\frac{p(x)}{q(x)} > m$. This set has probability $< \gamma/2$ under p. (ii) remaining numbers x such that $\log\frac{p(x)}{q(x)} < m$. This set has probability $\le \gamma/2^{m+1}$ under q, and therefore probability $\le \gamma/2$ under p. The lemma follows. □

VII-B. Weak compression and insurability

Unlike with strong compression, the connection of insurability with weak compression is not as clear cut. On the one hand, in Section VII-B.1 we first show two examples of distribution classes that are weakly compressible but not insurable. Then we follow it up with a distribution class $\mathcal{I}^\infty$ in Section VII-B.2 that is insurable but not weakly compressible.

VII-B.1) Weakly compressible but not insurable

We have already encountered two examples of distribution classes that are weakly compressible but not insurable.

For our first example, we consider the set $\mathcal{N}^\infty$ is the class of *i.i.d.* processes whose single letter marginals have finite moment. As shown in Example 3, $\mathcal{N}^\infty$ is weakly compressible but no neighborhood around any distribution $p \in \mathcal{N}$ is tight. Therefore, from Remark III-B that characterizes insurability from [13], $\mathcal{N}^\infty$ is not insurable.

For a second example, consider the collection $\mathcal{M}^\infty$ of all *i.i.d.* monotone sources. From Example 4, the class of monotone distributions is again weakly compressible, but no neighborhood around any distribution in $\mathcal{M}$ is tight. Therefore $\mathcal{M}^\infty$ is not insurable from Remark III-B.

VII-B.2) Insurable but not weakly compressible

In order to find *i.i.d.* measures that are insurable but not weakly compressible, we construct a class $\mathcal{I}$ of distributions over $\mathbb{N}$. As with other classes, $\mathcal{I}^\infty$ is the set of *i.i.d.* measures formed whose single letter marginals are $\mathcal{I}$.

To do so, first partition the set of natural numbers into the sets T_i, $i \ge 0$ as before. Recall again that

$$T_i = \{2^k, \ldots, 2^{k+1}-1\}.$$

Note that $|T_k| = 2^k$. Now, $\mathcal{I}$ is the collection of all possible distributions that can be formed as follows. For all $i \ge 1$, we pick exactly one element of T_i and assign it probability $1/(i(i+1))$. Note that $\mathcal{I}$ is not countable. Part of the rationale behind this construction is that for all $p \in \mathcal{I}$,

$$\sum_{n\ge 2^k} p(n) = \frac{1}{k+1},$$

namely, all tails are uniformly bounded over the class $\mathcal{I}$ to ensure insurability. Put another way, for all $\delta > 0$ and all distributions $p \in \mathcal{I}$,

$$F_p^{-1}(1-\delta) \le 2^{\lfloor \frac{1}{\delta} \rfloor} - 1.$$

We therefore have the following Corollary.

Corollary 9. The set $\mathcal{I}^\infty$ of measures is insurable. □

On the other hand,

Lemma 10. The set $\mathcal{I}^\infty$ of measures is not weakly compressible.

Proof Suppose q is any distribution over $\mathbb{N}$. We will show that $\exists p \in \mathcal{I}$ such that

$$\sum_{n\geq 1} p(n) \log \frac{p(n)}{q(n)}$$

is not finite. Since the entropy of every $p \in \mathcal{I}$ is finite, we just have to show that there $\exists p \in \mathcal{I}$ such that

$$\sum_{n\geq 1} p(n) \log \frac{1}{q(n)}.$$

Using the contrapositive of Corollary 2, we conclude that $\mathcal{I}^\infty$ is not weakly compressible.

Consider any distribution q over $\mathbb{N}$. Observe that for all i, $|T_i| = 2^i$. It follows that for all i there is $x_i \in T_i$ such that

$$q(x_i) \leq \frac{1}{2^i}.$$

But by construction, $\mathcal{I}$ contains a distribution p that has for its support $\{x_i : i \geq 1\}$ identified above. Furthermore p assigns

$$p(x_i) = \frac{1}{i(i+1)} \qquad \forall\, i \geq 1.$$

The KL divergence from p to q is not finite and the Lemma follows. □

VII-C. d.w.c. *and insurability*

We have already seen that *d.w.c.* implies insurability. The converse clearly does not hold—we have shown insurable classes that are not even weakly compressible, leave alone *d.w.c.*. We will however strengthen the observation by showing a weakly compressible, insurable class that is not *d.w.c.*.

Recall the collection $\mathcal{B}$ from Example 7. This class was shown to be not *d.w.c.*. However, $\mathcal{B}$ is tight and hence insurable. To see that $\mathcal{B}$ is also weakly compressible consider a distribution over $\mathbb{N}$ that for all $i \geq 0$, splits probability $1/i(i+1)$ equally among all 2^i elements of T_i. For all $p \in \mathcal{B}$, now

$$E_p \log \frac{1}{q(X)} < \infty$$

since every p has support at most 2. The class is then weakly compressible from [16].

VIII. Conclusion

We have identified statistical formulations that are positioned in between the notions of uniform and pointwise convergence. To place this in focus, we considered a universal compression problem that strengthens the notion of weak redundancy studied by [16]. This allows flexibility in modeling—we can allow rich model classes, yet figure out from the data when we are doing well. By introducing alternate prediction problems and comparing with it, we aim to highlight that the formulation of data-driven pointwise convergence is more general than just compression or prediction, and is one key to a deeper understanding of exactly what can be modeled from the data.

The formulation and analysis in this paper is only a start. Clearly even the compression framework has several open questions—can there be a guarantee on the excess codelength per sequence (instead of the strong redundancy that is an expectation over sequences)? This has to build on concentration results for the excess codelength when sources have countably infinite supports. Some initial work on these concentration results are in [].

More generally for other problems (in addition to compression and prediction), we conjecture that data derived convergence is effectively a way to have rich model classes that are locally simple. From a Bayesian perspective, we conjecture that data derived convergence is perhaps what is essential for model classes to allow for robust (as opposed to brittle []) priors.

As mentioned in the introduction, this paper can also be understood in terms of hierarchical universal coding which focusses more on whether the various subclasses are distinguishable. Furthermore, our formulation and analysis precludes any a-priori need to take a countable union of classes, but allows us to characterize arbitrary rich classes. However, we require that we provide data-driven consistent estimates no matter what the source. A natural question that is to be resolved then is, what if we have to provide the data-driven guarantees not for all sources but a set of sources with high probability under a given prior?

ACKNOWLEDGMENTS

N. Santhanam was supported by NSF Grants CCF-1065632 and CCF-1018984. V. Anantharam was supported by the ARO MURI grant W911NF- 08-1-0233, Tools for the Analysis and Design of Complex Multi-Scale Networks, the NSF grant CNS-0910702, the NSF Science & Technology Center grant CCF-0939370, Science of Information, Marvell Semiconductor Inc., and the U.C. Discovery program. A. Kavcic was supported by NSF Grant CCF-1018984. W. Szpankowski was supported by NSF Center for Science of Information (CSoI) Grant CCF-0939370, and in addition by NSA Grant H98230-11-1-0141, and NSF Grants DMS-0800568, and CCF-0830140.

REFERENCES

[1] M. Feder and N. Merhav. Hierarchical universal coding. *IEEE Transactions on Information Theory*, 42(5):1354–1364, 1996.
[2] C.M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
[3] V. Koltchinskii. Rademacher penalties and structural risk minimization. *IEEE Transactions on Information Theory*, 47:19021914, 2001.
[4] P.L. Bartlett, S. Boucheron, and G. Lugosi. Model selection and error estimation. *Machine Learning*, 48:85113, 2002.
[5] P.L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463482, 2002.
[6] Y.M. Shtarkov. Universal sequential coding of single messages. *Problems of Information Transmission*, 23(3):3—17, 1987.
[7] B. Fittingoff. Universal methods of coding for the case of unknown statistics. In *Proceedings of the 5th Symposium on Information Theory*, pages 129—135. Moscow-Gorky, 1972.
[8] M. Drmota and W. Szpankowski. Precise minimax redundancy and regrets. *IEEE Trans. Information Theory*, 50:2686–2707, 2004.
[9] Q. Xie and A. Barron. Minimax redundancy for the class of memoryless sources. *IEEE Transactions on Information Theory*, 43(2):647—657, March 1997.
[10] B.S. Clarke and A.R. Barron. Jeffreys' prior is asymptotically least favorable under entropy risk. *Journal of Statistical Planning and Inference*, 41(1):37—60, 1994.
[11] L.D. Davisson. Universal noiseless coding. *IEEE Transactions on Information Theory*, 19(6):783—795, November 1973.
[12] Narayana Santhanam. *Probability estimation and compression involving large alphabets*. PhD thesis, University of California, San Diego, 2006.
[13] N. Santhanam and V. Anantharam. Agnostic insurance of model classes. *Submitted to Journal of Machine Learning Research*, 2012. Full version available from arXiv doc id: 1212:3866.
[14] N. Merhav and M. Feder. Universal prediction. *IEEE Transactions on Information Theory*, 44(6):2124—2147, October 1998.
[15] M. Hosseini and N. Santhanam. On redundancy of memoryless sources over countable alphabets. Submitted for publication. Full version available from arXiv.
[16] J.C. Kieffer. A unified approach to weak universal source coding. *IEEE Transactions on Information Theory*, 24(6):674—682, November 1978.
[17] J. Dugundji. *Topology*. Allyn and Bacon Inc., Boston, 1970.
[18] A. Orlitsky, N.P. Santhanam, and J. Zhang. Universal compression of memoryless sources over unknown alphabets. *IEEE Transactions on Information Theory*, 50(7):1469—1481, July 2004.
[19] A. Orlitsky, N.P. Santhanam, K. Viswanathan, and J.Zhang. On modeling profiles instead of values. In *Uncertainty in Artificial Intelligence*, 2004.
[20] K.L. Chung. A note on the ergodic theorem of information theory. *Annals of Mathematical Statistics*, 32:612—614, 1961.
[21] S. Ho and R. Yeung. On information divergence measures and joint typicality. *IEEE Transactions on Information Theory*, 56(12):58935905, 2010.

APPENDIX I
INSURABILITY AND *d.w.c.* CONSISTENT ESTIMATION OF PERCENTILES

The proof that consistent upper bounding of all percentiles of a distribution in a *d.w.c.* manner is identical to insurability follows from a careful reading of [13]. This section mainly highlights those relevant results from [13].

Suppose we have a *d.w.c.* consistent bounds of all percentiles of a distribution. We can then construct an insurance scheme by setting as premiums the *d.w.c.* obtained bounds on the $1 - \frac{1}{n^2}$ percentile, where n is the length of the candidate string on which the *d.w.c.* scheme yields bounds. Looking at the contrapositive, if no insurance scheme is possible, there can be no *d.w.c.* consistent estimate either. Hence, local tightness of a model class is necessary for *d.w.c.* consistent estimate of all percentiles.

That the local tightness of a model class is sufficient follows from a more careful reading of [13]. The sufficiency proof for insurability (shown in [13]) covers the model class with a countable number of neighborhoods similar to the quantization in Section VI-B, where each neighborhood is locally tight (instead of the strong compressibility condition used in this paper).

Note that in any locally tight class, we can trivially upper bound all percentiles of a distribution in that class. The only question then is whether we can identify the locally tight neighborhood the underlying distribution belongs

to from the data. Namely, using notation of this paper, can we ensure that we fall into good traps? More formally, given any $\eta > 0$, can we make the probability of falling into bad traps $< \eta$? We show in [13] that the probability of ever falling into a bad trap can be made arbitrarily small. Therefore, if a model class is locally tight, we can obtain consistent upper bounds on all percentiles of the underlying distribution in a *d.w.c.* manner.

APPENDIX II
TECHNICAL LEMMAS

The proofs of Lemmas 11 and 12 be found in [13].

Lemma 11. Let p and q be probability distributions on $\mathbb{N}$. Then

$$\frac{1}{4\ln 2}|p-q|_1^2 \le \mathcal{J}(p,q) \le \frac{1}{\ln 2}|p-q|_1 \ .$$

If, in addition, r is a probability distribution on $\mathbb{N}$, then

$$\mathcal{J}(p,q) + \mathcal{J}(q,r) \ge \mathcal{J}^2(p,r)\frac{\ln 2}{8}. \qquad \square$$

Lemma 12. Let p and q be probability distributions on a countable set $\mathcal{A}$ with $\mathcal{J}(p,q) \le \epsilon$. Let p^N and q^N be distributions over $\mathcal{A}^N$ obtained by *i.i.d.* sampling from p and q respectively (the distribution induced by the product measure). For any $R_N \subset \mathcal{A}^N$ and $\alpha > 0$, if $p^N(R_N) \ge 1-\alpha$, then

$$q^N(R_N) \ge 1 - \alpha - 2N^3\sqrt{4\epsilon \ln 2} - \frac{1}{N}. \qquad \square$$

APPENDIX III
BOUND ON $D_n(p||q)$ USING SINGLE LETTER CONSTRAINTS ON p AND q

We reproduce the proof here that we can construct a measure q^* such that for $p \in \mathcal{M}_h$,

$$\sup_{p\in\mathcal{M}_h^n} \frac{1}{n} D_n(p||q) \le \frac{2h^{3/2}}{\sqrt{\log n}} + \pi\sqrt{\frac{2}{3n}}$$

This result is borrowed from [] which considers in detail the connections between single letter and length-n average redundancy. Moreover, the paper also shows that we could have the single letter average redundancy be finite, but the normalized length n redundancy is bounded away from 0 in the limit as $n \to \infty$.

To see this, consider the distribution over $\mathbb{N}$ $q(i) = 1/i(i+1)$. As shown in Example 4, we have

$$\sup_{p\in\mathcal{M}_h} E_p\left(\log\frac{1}{q(X)}\right)^2 < 4h.$$

Consider a scheme that encodes patterns [18] of symbols first followed by an encoding using $\log\frac{1}{q(x)}$ bits to describe every symbol x that appeared in the string. The expected (not normalized by n) additional number of bits to encode the pattern is $\pi\sqrt{\frac{2}{3}n}$ using the results in [18], while the expected number of bits to describe the symbols of length-n strings is

$$\sum_{i\in\mathbb{N}} (1-(1-p_i)^n)\log\frac{1}{q(i)}.$$

Therefore, the expected extra bits the scheme uses for length-n strings is (without normalizing by n) $\pi\sqrt{\frac{2}{3}n}$ plus

$$\begin{aligned}
&\le \sum_{i\in\mathbb{N}} (1-(1-p_i)^n)\log\frac{1}{q(i)} \\
&\overset{(a)}{\le} \sqrt{\sum_{i\in\mathbb{N}} (1-(1-p_i)^n)\sum_{j\in\mathbb{N}} (1-(1-p_j)^n)\left(\log\frac{1}{q(j)}\right)^2} \\
&\le \sqrt{(\mathbb{E}M_n)(4nh)} \\
&\overset{(b)}{\le} \frac{2nh^{3/2}}{\sqrt{\log n}},
\end{aligned}$$

where (a) follows from the Cauchy-Schwarz inequality, and (b) from [19, Lemma ?].

APPENDIX IV
Φ ENTERS WITH PROBABILITY 1

We reproduce the argument from [13] here for completeness.

Every distribution $p \in \mathcal{P}$ is contained in at least one of the elements of the cover $\mathcal{Q}_{\tilde{\mathcal{P}}}$. Recall the enumeration of $\tilde{\mathcal{P}}$. Let p' be be centroid with the smallest index among all centroids in $\tilde{\mathcal{P}}$ whose zones contain p. Let Q be the zone of p'. There is thus some $\gamma > 0$ such that the neighborhood around p given by

$$I(p, \gamma) := \{\tau : |p - \tau|_1 < \gamma\}$$

satisfies $I(p, \gamma) \subseteq Q$. Note in particular that p is in the reach of p'. With probability 1, sequences generated by p will have their empirical distribution within $I(p, \gamma)$ (see [20] or Lemma 14 for an alternate proof). Next (5) will hold for all sequences whose empirical distributions that fall in $I(p, \gamma)$ whose length n is large enough—since $C(p')$ and $\iota(p')$ do not change with n, the right hand side diminishes to zero polynomially with n while the left hand side diminishes exponentially to zero. Therefore, with probability 1, sequences generated by p are trapped. Next, (6) will also hold almost surely, for if τ is the empirical probability of sequences generated by p, then (with a little abuse of notation)

$$F_\tau^{-1}(1 - \sqrt{D_{p'}}/6) \to F_p^{-1}(1 - \sqrt{D_{p'}}/6)$$

with probability 1. Note that the quantity on the left is actually a random variable that is sequence dependent (since τ is the empirical distribution of the sequence). Furthermore, we also have

$$\begin{aligned} 2F_p^{-1}(1 - \sqrt{D_{p'}}/6) &\leq 2\left(\sup_{r \in B(p', \epsilon_{p'})} F_r^{-1}(1 - \sqrt{D_{p'}}/6)\right) \\ &= \log C(p'), \end{aligned}$$

where the first inequality follows since p is in the reach of p'.

APPENDIX V
PROBABILITY OF FALLING INTO BAD TRAPS

We need two additional technical lemmas for this proof. The proofs of both are in [13].

Lemma 13. Let $\epsilon_0 > 0$. If

$$|p_0 - q|_1 \leq \frac{\epsilon_0^2 (\ln 2)^2}{16},$$

then for all $p \in \mathcal{P}$ with $\mathcal{J}(p, p_0) \geq \epsilon_0$, we have

$$\mathcal{J}(p, q) \geq \frac{\epsilon_0^2 \ln 2}{16}.$$ □

Lemma 14. Let p be any probability distribution on $\mathbb{N}$. Let $\delta > 0$ and let $k \geq 2$ be an integer. Let X_1^n be a sequence generated *i.i.d.* with marginals p and let $q(X^n)$ be the empirical distribution of X_1^n. Then

$$\begin{aligned} &p\big(|q(X^n) - p| > \delta \text{ and } 2F_q^{-1}(1 - \delta/6) \leq k\big) \\ &\leq (2^k - 2) \exp\left(-\frac{n\delta^2}{18}\right). \end{aligned}$$

Remark There is a lemma that looks somewhat similar in [21]. The difference from [21] is that the right side of the inequality above does *not* depend on p, and this property is crucial for its use here. □

Let τ be any length-n empirical distribution trapped by $\tilde{p}$ with reach $\tilde{\epsilon}$ such that $p \notin B(\tilde{p}, \epsilon_{\tilde{p}})$. If p is "far" from $\tilde{p}$ (because p is not in $\tilde{p}$'s reach), namely

$$\mathcal{J}(\tilde{p}, p) \geq \tilde{\epsilon},$$

but τ is "close" to $\tilde{p}$ (because τ has to be in $\tilde{p}$'s zone to be captured by it), namely

$$|\tilde{p} - \tau|_1 < \frac{\tilde{\epsilon}^2 (\ln 2)^2}{16},$$

then we would like τ to be far from p. That is exactly what we obtain from the triangle-inequality like Lemma 13, namely that

$$\mathcal{J}(p,\tau) \geq \frac{\tilde{\epsilon}^2 \ln 2}{16}$$

and hence, for all τ trapped by $\tilde{p}$ that

$$|p-\tau|_1^2 \geq \mathcal{J}^2(p,\tau)(\ln 2)^2 \geq \frac{\tilde{\epsilon}^4 (\ln 2)^4}{256} = D_{\tilde{p}}^2.$$

We need not be concerned that the right side above depends on $\tilde{p}$, and there may be actually no way to lower bound the rhs as a function of just p. Rather, we take care of this issue by setting the entry point appropriately via (5). Thus, for $p \in \mathcal{P}^\infty$, the probability length-n sequences with empirical distribution τ is trapped by a bad $\tilde{p}$ is, using (5) and (6)

$$\begin{aligned}
&\leq p\left(|\tau - p|^2 \geq D_{\tilde{p}} \text{ and } 2F_\tau^{-1}(1 - \frac{\sqrt{D_{\tilde{p}}}}{6}) \leq \log C(\tilde{p}) \right) \\
&= p\left(|\tau - p| \geq \sqrt{D_{\tilde{p}}} \text{ and } F_\tau^{-1}(1 - \sqrt{D_{\tilde{p}}}/3) \leq \log C(\tilde{p}) \right) \\
&\overset{(a)}{\leq} (C(\tilde{p}) - 2) \exp\left(-\frac{n D_{\tilde{p}}}{18} \right) \\
&\overset{(b)}{\leq} \frac{\eta (C(\tilde{p}) - 2)}{2C(\tilde{p}) \iota(\tilde{p})^2 n(n+1)} \\
&\leq \frac{\eta}{2\iota(\tilde{p})^2 n(n+1)},
\end{aligned}$$

where the inequality (a) follows from Lemma 14 and (b) from (5). Therefore, the probability of sequences falling into bad traps

$$\leq \sum_{n \geq 1} \sum_{\tilde{p} \in \tilde{\mathcal{P}}} \frac{\eta}{2\iota(\tilde{p})^2 n(n+1)} \leq \eta/2$$

since $\sum_{\tilde{p} \in \tilde{\mathcal{P}}} \frac{1}{\iota(\tilde{p})^2} \leq \sum_{n \geq 1} \frac{1}{n(n+1)} = 1$.